Journal of Machine Learning Research 1 (2015) 1-48 Submitted 9/15; Published [[pub date]]

Recurrent Exponential-Family Harmoniums without Backprop-through-Time

Joseph G. Makin MAKIN@PHY.UCSF.EDU
Benjamin K. Dichter BEN.DICHTER@BERKELEY.EDU
Philip N. Sabes SABES@PHY.UCSF.EDU
Center for Integrative Neuroscience
University of California, San Francisco
San Francisco, CA 94143-0444, USA

Editor: [[editor's name]]

arXiv:1605.05799v1 [cs.LG] 19 May 2016

Abstract

Exponential-family harmoniums (EFHs), which extend restricted Boltzmann machines (RBMs) from Bernoulli random variables to other exponential families (Welling et al., 2005), are generative models that can be trained with unsupervised-learning techniques, like contrastive divergence (Hinton et al., 2006; Hinton, 2002), as density estimators for static data. Methods for extending RBMs—and likewise EFHs—to data with temporal dependencies have been proposed previously (Sutskever and Hinton, 2007; Sutskever et al., 2009), the learning procedure being validated by qualitative assessment of the generative model. Here we propose and justify, from a very different perspective, an alternative training procedure, proving sufficient conditions for optimal inference under that procedure. The resulting algorithm can be learned with only forward passes through the data—backprop-through-time is not required, as in previous approaches. The proof exploits a recent result about information retention in density estimators (Makin and Sabes, 2015), and applies it to a "recurrent EFH" (rEFH) by induction. Finally, we demonstrate optimality by simulation, testing the rEFH: (1) as a filter on training data generated with a linear dynamical system, the position of which is noisily reported by a population of "neurons" with Poisson-distributed spike counts; and (2) with the qualitative experiments proposed by Sutskever et al. (2009).

Keywords: RBM, Time-Series Data, Unsupervised Learning, Exponential Family Harmonium, identifiable

1. Introduction

We are interested in unsupervised learning of data with temporal dependencies. The aim of such learning is to extract the statistical structure from "trajectories" or sequences, after which it will be possible, for example, to predict the next frame in the sequence, to make optimal inferences about the dynamical state of the underlying system being observed ("filter"), or to generate novel sequences with the same statistical properties.

Ideally, we would make use of some of the sophisticated unsupervised-learning techniques developed for static data. If, however, a single training vector consists of the observations at a single time step, the temporal dependencies obviously will not be learned. The opposite extreme, feasible in some contexts, is to assign to a single training vector an entire trajectory

©2015 Joseph G. Makin, Benjamin K. Dichter, and Philip N. Sabes.

of observations; but in general, trajectories may be very long, vary in length, or required to be learned on-line.

A middle course is to try to construct, at each time step, *augmented* observations that are shorter than an entire trajectory of observations, but that form, unlike the raw observations, a Markov chain. When the *direct* temporal dependencies are not, indeed, the length of the trajectory itself, past observations' bearing on future observations can be summarized compactly. Data consisting of the current observation, augmented by such a summary, form a Markov chain across time, which facilitates the application of recursive procedures for filtering and prediction. In the most familiar example of this strategy, the Kalman filter, the past noisy observations of a discrete-time, linear-Gaussian dynamical system are summarized by the posterior cumulants over a hidden dynamical state. These cumulants, which are also the sufficient statistics for the hidden state (conceived as a parameter), are the state-estimate vector, whose length is proportional to the length of the direct temporal dependencies, and a covariance matrix that encodes the (inverse) reliability of this state estimate. The recursive filtering procedure in turn makes parameter learning possible—in this case, with an expectation-maximization (EM) algorithm (see, e.g., Ghahramani and Hinton, 1996).

Our goal is to extend this "middle" strategy to more general temporal sequences, for which we are unable, for example, to compute analytically the posterior distribution over the hidden state, or even to formulate equations of state or an observation model. The compact summaries, consequently, we shall require the algorithm itself to formulate, rather than deriving them analytically. In particular, our basic building block will be the exponential-family harmonium (EFH) (Welling et al., 2005), a two-layer Markov random field with full interlayer connections and no intralayer connections, in which the conditional distributions over the hidden layer given the visible layer and vice versa are products over exponential-family distributions. (The restricted Boltzmann machine corresponds to the special case of Bernoulli conditional distributions.) For the summaries, with which we augment the observations at each time step, we choose the hidden-unit means from the previous time step.

We have recently (Makin et al., 2015) introduced this temporal extension to the EFH, the "recurrent EFH" (rEFH); investigated its predictions for receptive-field formation in primate cerebral cortex; and compared its performance as a filter for simple dynamical systems against Kalman filters learned with EM. Here, we investigate the analytical properties of the rEFH, supplying sufficient conditions and a proof for when the training procedure results in a model that provides optimal inference. The central idea behind the proof is to abandon the assumption of model correctness that underlies use of, e.g., linear-Gaussian dynamical systems, in favor of a weaker assumption about model identifiability. This assumption, which holds for the rEFH, guarantees that, at each time step, the sufficient statistics for the hidden units of the trained model are also sufficient statistics for the current *true* hidden state. That is, the model's hidden units learn to encode the compact summary corresponding to the filter cumulants.

The rEFH is very similar to, but importantly different from, two recent temporal extensions to the EFH, the TRBM (Sutskever and Hinton, 2007) and the RTRBM (Sutskever et al., 2009). We show why the proof does not apply to these models, and then compare all three empirically by training and testing on data where the optimal posterior estimate of

the true dynamical state is computable analytically. We also compare model performance on the synthetic data set used in those papers: the rEFH is superior to the TRBM and competitive with the RTRBM, despite a simpler learning algorithm. Finally, we show that generating sample trajectories—*in the reverse direction*—is also simple in the rEFH, in contrast to the TRBM and RTRBM, for which generation is hard.

2. Theory

We begin by describing our strategy for extending the EFH to temporal data (Section 2.1) before justifying it more rigorously (Section 2.2), and finally describing the training procedure in detail (Section 2.3). Our aim is a training procedure and architecture that can be guaranteed to provide *optimal inference*. That is, after training, our model should allow us to compute the posterior distribution, $p(\mathbf{x}_t|\mathbf{y}_0, \ldots, \mathbf{y}_t)$, over the dynamical state, $\mathbf{X}_t$, of the underlying causes of our observations, $\mathbf{y}_t$, conditioned on the preceding sequence of such observations. In particular, we require that the (mean) hidden-unit activities of the model at time t encode the cumulants of this distribution.

2.1 Strategy for extending the EFH to temporal data

We consider temporal data $\mathbf{Y}_0, \ldots, \mathbf{Y}_T$ that have been generated by a very general discrete-time dynamical process, anything consistent with the independence statements asserted by the hidden Markov model in the top two rows of Fig. 1a, without any other assumptions about the distributions parameterizing it. (The Markov assumption is weak in the sense that longer temporal dependencies can be accommodated under it simply by lengthening the state vector.) We shall, however, introduce below some constraints relating to a technical condition called "identifiability." Our goal is to make optimal inferences about the underlying "causes," $\mathbf{X}_0, \ldots, \mathbf{X}_T$, for the observations, $\mathbf{Y}_0, \ldots, \mathbf{Y}_T$, with joint distribution

$$p(\mathbf{x}_0, \ldots, \mathbf{x}_T, \mathbf{y}_0, \ldots, \mathbf{y}_T) = \prod_{t=0}^{T} p(\mathbf{y}_t|\mathbf{x}_t)p(\mathbf{x}_t|\mathbf{x}_{t-1}),$$

corresponding to Fig. 1a. (Here as elsewhere, for notational compactness, $p(\mathbf{x}_0|\mathbf{x}_{-1}) := p(\mathbf{x}_0)$, and likewise for other expressions that reference negative times.) Given nothing more than samples from a "data distribution" of this general form, we construct our own latent-variable model, q, and attempt to bring the probabilities it assigns to observation sequences close to the probabilities with which they occur in the data; i.e., we seek parameters $\boldsymbol{\theta}$ such that $q(\mathbf{y}_0, \ldots, \mathbf{y}_T; \boldsymbol{\theta}) = p(\mathbf{y}_0, \ldots, \mathbf{y}_T)$ ("density estimation").

Here we introduce an important distinction. We do not assume that the latent variables of the model have the same form as those "in the world," $\mathbf{X}$, so in fact we use a different letter for them, $\mathbf{Z}$. In training latent-variable density estimators, the correspondence between $\mathbf{Z}$ and $\mathbf{X}$ is not always straightforward, even after the model has become a good one for the observed data. Nevertheless, our recent result about the character of this correspondence (Makin and Sabes, 2015) can be used to show how to achieve our goal, optimal inference of the true causes, $\mathbf{X}$.

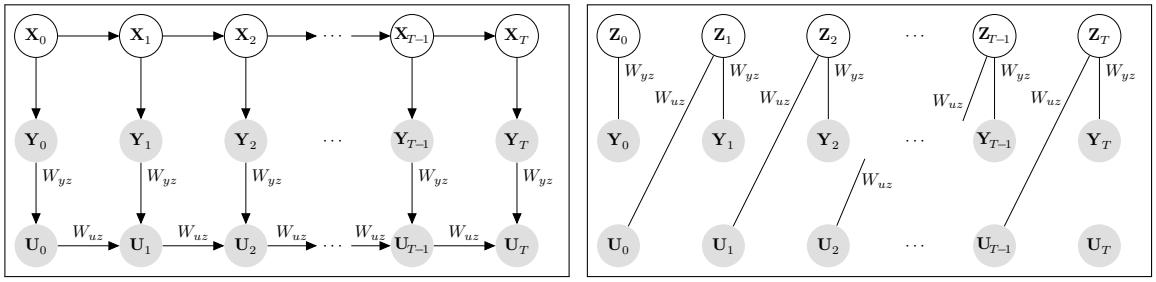
$\mathbf{X}_0$ $\mathbf{X}_1$ $\mathbf{X}_2$ $\cdots$ $\mathbf{X}_{T-1}$ $\mathbf{X}_T$
$\mathbf{Y}_0$ $\mathbf{Y}_1$ $\mathbf{Y}_2$ $\cdots$ $\mathbf{Y}_{T-1}$ $\mathbf{Y}_T$
W_{yz}
$\mathbf{U}_0$ $\mathbf{U}_1$ $\mathbf{U}_2$ $\cdots$ $\mathbf{U}_{T-1}$ $\mathbf{U}_T$
W_{uz}
$\mathbf{Z}_0$ $\mathbf{Z}_1$ $\mathbf{Z}_2$ $\cdots$ $\mathbf{Z}_{T-1}$ $\mathbf{Z}_T$
W_{yz}
W_{uz}
$\mathbf{Y}_0$ $\mathbf{Y}_1$ $\mathbf{Y}_2$ $\cdots$ $\mathbf{Y}_{T-1}$ $\mathbf{Y}_T$
$\mathbf{U}_0$ $\mathbf{U}_1$ $\mathbf{U}_2$ $\cdots$ $\mathbf{U}_{T-1}$ $\mathbf{U}_T$

(a) (b)

Figure 1: Two generative models for the observation sequences. (a) The data distribution. The bottom row is computed deterministically, $\mathbf{U}_t := \mathbb{E}_q[\mathbf{Z}_t|\mathbf{y}_t, \mathbf{u}_{t-1}; \boldsymbol{\theta}]$. (b) The rEFH. At each time step t, an EFH with hidden layer $\mathbf{Z}_t$ is trained on "observations" $\mathbf{U}_{t-1}$ and $\mathbf{Y}_t$. The units in each layer are conditionally independent given the other layer, but the layers have been collapsed here to single nodes to save space (cf. Fig. 2b). Weight matrices are labeled; biases are not shown.

Consider in particular observed (static) random variables $\mathbf{V}$. In many cases one can say that the sufficient statistics[1] for the model's latent variables, $\mathbf{U}_\mathbf{Z}(\mathbf{V})$, are also sufficient statistics for the true latent variables, $\mathbf{X}$. This is intuitive: the same aspects of $\mathbf{V}$ are in need of explanation, whether via the world or via the model. The result does not hold for all latent-variable density estimators, but for a class of such models that obeys a set of rather weak technical conditions (see Makin and Sabes, 2015, and Theorem 2 below), it is the case that

$$\forall \mathbf{v} \quad q(\mathbf{v}; \boldsymbol{\theta}) = p(\mathbf{v}) \implies \mathcal{I}(\mathbf{U}_\mathbf{Z}(\mathbf{V}); \mathbf{X}) = \mathcal{I}(\mathbf{V}; \mathbf{X}), \tag{1}$$

with $\mathcal{I}(\cdot;\cdot)$ the mutual information. That is, after a good model has been learned (left-hand equation), transforming $\mathbf{V}$ into any set of sufficient statistics for $\mathbf{Z}$ discards no information about the underlying cause, $\mathbf{X}$. When information about $\mathbf{X}$ reaches $\mathbf{Z}$ *only* via $\mathbf{V}$, Eq. 1 is equivalent (see Makin and Sabes, 2015) to

$$\forall \mathbf{v} \quad q(\mathbf{v}; \boldsymbol{\theta}) = p(\mathbf{v}) \implies \forall \mathbf{v}, \mathbf{x} \quad p(\mathbf{x}|\mathbf{u}_\mathbf{Z}(\mathbf{v})) = p(\mathbf{x}|\mathbf{v}). \tag{2}$$

In words, transforming $\mathbf{V}$ into sufficient statistics for $\mathbf{Z}$ does not change the posterior distribution over the underlying cause, $\mathbf{X}$. For EFHs, the sufficient statistics $\mathbf{U}_\mathbf{Z}(\mathbf{V})$ are the posterior mean, $\mathbb{E}_q[\mathbf{Z}|\mathbf{V}]$; hence, in a trained model, the mean activity of the hidden units captures all the information about the true cause, $\mathbf{X}$, as was contained in the vector of observed responses, $\mathbf{V}$.

We construct a procedure for training the EFH on data with temporal dependencies that exploits this result. In particular, at every time step, we train the EFH on an "augmented" input vector $\mathbf{v}_t$ that consists of the current observations, $\mathbf{y}_t$, concatenated together with the sufficient statistics for $\mathbf{Z}_{t-1}$ from the previous time step, $\mathbf{u}_{\mathbf{Z}_{t-1}}(\mathbf{v}_{t-1})$—i.e., the previous posterior mean, $\mathbb{E}_q[\mathbf{Z}_{t-1}|\mathbf{v}_{t-1}]$. The central intuition for why this procedure will produce models providing optimal inference is as follows. Training the model at the first time

1. "T," the usual symbol for sufficient statistics, is already used for time, so we use "U" instead.

step ($t = 0$) forces the hidden units—or, more precisely, the sufficient statistics for them, $\mathbb{E}_q[\mathbf{Z}_0|\mathbf{y}_0]$—to extract from observations $\mathbf{y}_0$ a good summary of the underlying state, $\mathbf{X}_0$, no information about it being lost. The sufficient statistics at the second time step ($t = 1$), then, are forced to become a good "meta-summary" of both the information in $\mathbf{Y}_1$ about $\mathbf{X}_1$ *and* the previous summary. Continuing in this fashion builds up in these sufficient statistics—the means of the hidden units—a good running summary about the true state and its history. We call the resulting model the recurrent, exponential-family harmonium (rEFH).

2.2 Rigorous justification of the training procedure

We now formalize this intuitive argument. We begin with the result from our previous work:

Theorem 1 *Let* $p(\mathbf{v}) = \sum_{\mathbf{x}} p(\mathbf{v}|\mathbf{x})p(\mathbf{x})$ *and* $q(\mathbf{v};\boldsymbol{\theta}) = \sum_{\mathbf{z}} p(\mathbf{v}|\mathbf{z};\boldsymbol{\theta})p(\mathbf{z};\boldsymbol{\theta})$ *be expressible as* identifiable *finite mixtures of the same conditional distribution* $f(\mathbf{v}|\boldsymbol{\lambda})$, $\sum_{\boldsymbol{\lambda}} f(\mathbf{v}|\boldsymbol{\lambda})p(\boldsymbol{\lambda})$ *and* $\sum_{\boldsymbol{\lambda}} f(\mathbf{v}|\boldsymbol{\lambda})q(\boldsymbol{\lambda})$*, respectively. That is, the map from prior distributions over "parameters"* $\boldsymbol{\Lambda}$ *to mixtures (marginal distributions over* $\mathbf{V}$*) is injective. Then Eq. 1 holds. Furthermore, if* $\mathbf{X}$ *and* $\mathbf{Z}$ *are independent conditioned on* $\mathbf{V}$*, then Eq. 2 holds as well.*

For the proof, see the work of Makin and Sabes (2015). Here we prove by induction that the training procedure just proposed generalizes this theorem to temporal data.

Theorem 2 *For all* t*, let the marginal distributions*

$$p(\mathbf{v}_t;\boldsymbol{\theta}) := \sum_{\mathbf{x}_0,\ldots,\mathbf{x}_t;\boldsymbol{\theta}} p(\mathbf{v}_t|\mathbf{x}_0,\ldots,\mathbf{x}_t)p(\mathbf{x}_0,\ldots,\mathbf{x}_t)$$
$$q(\mathbf{v}_t;\boldsymbol{\theta}) := \sum_{\mathbf{z}_t} q(\mathbf{v}_t|\mathbf{z}_t;\boldsymbol{\theta})q(\mathbf{z}_t;\boldsymbol{\theta})$$

be expressible as identifiable *finite mixtures of the same conditional distribution* $f(\mathbf{v}_t|\boldsymbol{\lambda}_t)$, $\sum_{\boldsymbol{\lambda}_t} f(\mathbf{v}_t|\boldsymbol{\lambda}_t)p(\boldsymbol{\lambda}_t)$ *and* $\sum_{\boldsymbol{\lambda}_t} f(\mathbf{v}_t|\boldsymbol{\lambda}_t)q(\boldsymbol{\lambda}_t)$*, respectively. Furthermore, let the distribution of observed data* $\mathbf{V}_t$ *be consistent with the independence statements asserted by Fig. 1a, in which in particular each observation consists of the emissions of a hidden Markov model, concatenated together with the sufficient statistics,* $\mathbf{U}_{\mathbf{Z}_{t-1}}(\mathbf{V}_{t-1})$*, for the latent variables* $\mathbf{Z}_{t-1}$ *of* q *at the preceding time step:*

$$\mathbf{V}_t := [\mathbf{U}_{\mathbf{Z}_{t-1}}(\mathbf{V}_{t-1}), \mathbf{Y}_t]. \tag{3}$$

Then

$$\forall t, \forall \mathbf{v}_t \;\; q(\mathbf{v}_t;\boldsymbol{\theta}) = p(\mathbf{v}_t;\boldsymbol{\theta}) \implies \mathcal{I}(\mathbf{U}_{\mathbf{Z}_t}(\mathbf{V}_t);\mathbf{X}_t) = \mathcal{I}(\mathbf{Y}_0,\ldots,\mathbf{Y}_t;\mathbf{X}_t). \tag{4}$$

If, moreover, $\mathbf{X}_t$ *and* $\mathbf{Z}_t$ *are independent conditioned on* $\mathbf{V}_t$*, then*

$$\forall t, \forall \mathbf{v}_t \;\; q(\mathbf{v}_t;\boldsymbol{\theta}) = p(\mathbf{v}_t;\boldsymbol{\theta}) \implies p(\mathbf{x}_t|\mathbf{u}_{\mathbf{Z}_t}(\mathbf{v}_t)) = p(\mathbf{x}_t|\mathbf{y}_0,\ldots,\mathbf{y}_t). \tag{5}$$

Note that the variables $\mathbf{u}_{\mathbf{Z}_{t-1}}(\mathbf{v}_{t-1})$ can be calculated deterministically from $\mathbf{y}_0,\ldots,\mathbf{y}_{t-1}$, so the data distribution p induces a distribution over these sufficient statistics as well as

the observations. This is the joint distribution we have called $p(\mathbf{v}_t;\boldsymbol{\theta})$, whose generative model is depicted in Fig. 1a. The equality on the left of Eq. 4 says that the EFH is a good density estimator for its inputs at time t, $\mathbf{V}_t$ (cf. Fig. 1b). The information equality on the right says that all the information about the current state, $\mathbf{X}_t$, that was available in all the observations up to the present, $\mathbf{y}_0,\ldots,\mathbf{y}_t$, is also available in a particular function, $\mathbf{u}_{\mathbf{Z}_t}(\cdot)$, of just the current augmented input vector, $\mathbf{v}_t$. Just as Eq. 1 can be re-expressed as Eq. 2, this information equality can be alternatively expressed as Eq. 5.

The critical technical condition guaranteeing Eq. 1, on which the proof depends, is that $p(\mathbf{v}_t)$ and $q(\mathbf{v}_t)$ be expressible in terms of the same finite-mixture model, and that the map from distributions over parameters of this mixture to the marginal distribution over $\mathbf{V}_t$ be injective, except possibly on sets of measure zero. Precise characterization being rather technical, we defer it to the **Appendix A**. We now prove theorem Theorem 2 by induction, under the assumption that these technical conditions hold.

Proof It is easily verified that $\mathbf{X}_t$ and $\mathbf{Z}_t$ are independent conditioned on $\mathbf{V}_t$, for all time, so we shall work throughout with Eqs. 2 and 5 rather than Eqs. 1 and 4. Now, the base case of the induction is simply an application of Eq. 2. Assume that $q(\mathbf{v}_0;\boldsymbol{\theta}) = p(\mathbf{v}_0)$, or equivalently that $q(\mathbf{y}_0;\boldsymbol{\theta}) = p(\mathbf{y}_0)$ (since at $t=0$ there is no recurrent vector to concatenate with the current observation $\mathbf{y}_0$; see Fig. 1b); that is, that the EFH has become a good density estimator for the initial data. Then Eq. 2 implies that

$$p(\mathbf{x}_0|\mathbf{u}_{\mathbf{Z}_0}(\mathbf{v}_0)) = p(\mathbf{x}_0|\mathbf{u}_{\mathbf{Z}_0}(\mathbf{y}_0)) = p(\mathbf{x}_0|\mathbf{y}_0). \tag{6}$$

Now for the induction step, assume that

$$p(\mathbf{x}_{t-1}|\mathbf{u}_{\mathbf{Z}_{t-1}}(\mathbf{v}_{t-1})) = p(\mathbf{x}_{t-1}|\mathbf{y}_0,\ldots,\mathbf{y}_{t-1}). \tag{7}$$

We want to show that under this assumption, the EFH having become a good density estimator for its inputs at time t, $q(\mathbf{v}_t;\boldsymbol{\theta}) = p(\mathbf{v}_t;\boldsymbol{\theta})$, implies

$$p(\mathbf{x}_t|\mathbf{u}_{\mathbf{Z}_t}(\mathbf{v}_t)) = p(\mathbf{x}_t|\mathbf{y}_0,\ldots,\mathbf{y}_t). \tag{8}$$

Together with the base case, this will show that the implication holds for all time—Eq. 5.

To save space we abbreviate $\mathbf{U}_t := \mathbf{U}_{\mathbf{Z}_t}(\mathbf{V}_t)$ and omit references to the parameters. Now, since $q(\mathbf{v}_t) = p(\mathbf{v}_t)$, Eq. 2 implies that

$$\begin{aligned}
p(\mathbf{x}_t|\mathbf{u}_t) &= p(\mathbf{x}_t|\mathbf{v}_t) \\
&= p(\mathbf{x}_t|\mathbf{y}_t,\mathbf{u}_{t-1}) \\
\implies p(\mathbf{x}_t|\mathbf{u}_t)p(\mathbf{y}_t|\mathbf{u}_{t-1}) &= p(\mathbf{x}_t|\mathbf{y}_t,\mathbf{u}_{t-1})p(\mathbf{y}_t|\mathbf{u}_{t-1}) && \text{(9a)} \\
&= p(\mathbf{y}_t|\mathbf{x}_t,\mathbf{u}_{t-1})p(\mathbf{x}_t|\mathbf{u}_{t-1}) \\
&= p(\mathbf{y}_t|\mathbf{x}_t)\sum_{\mathbf{x}_{t-1}} p(\mathbf{x}_t,\mathbf{x}_{t-1}|\mathbf{u}_{t-1}) && \text{(9b)} \\
&= p(\mathbf{y}_t|\mathbf{x}_t)\sum_{\mathbf{x}_{t-1}} p(\mathbf{x}_t|\mathbf{x}_{t-1},\mathbf{u}_{t-1})p(\mathbf{x}_{t-1}|\mathbf{u}_{t-1}) \\
&= p(\mathbf{y}_t|\mathbf{x}_t)\sum_{\mathbf{x}_{t-1}} p(\mathbf{x}_t|\mathbf{x}_{t-1})p(\mathbf{x}_{t-1}|\mathbf{u}_{t-1}) && \text{(9c)}
\end{aligned}$$

$$= p(\mathbf{y}_t|\mathbf{x}_t) \sum_{\mathbf{x}_{t-1}} p(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{y}_0, \ldots, \mathbf{y}_{t-1}) p(\mathbf{x}_{t-1}|\mathbf{y}_0, \ldots, \mathbf{y}_{t-1}) \quad (9d)$$

$$= p(\mathbf{y}_t|\mathbf{x}_t) p(\mathbf{x}_t|\mathbf{y}_0, \ldots, \mathbf{y}_{t-1})$$

$$= p(\mathbf{y}_t|\mathbf{x}_t, \mathbf{y}_0, \ldots, \mathbf{y}_{t-1}) p(\mathbf{x}_t|\mathbf{y}_0, \ldots, \mathbf{y}_{t-1}) \quad (9e)$$

$$= p(\mathbf{x}_t|\mathbf{y}_0, \ldots, \mathbf{y}_t) p(\mathbf{y}_t|\mathbf{y}_0, \ldots, \mathbf{y}_{t-1}). \quad (9f)$$

The independence conditions used in the transitions to Eqs. 9b, 9c, 9d, and 9e are licensed by the graph in Fig. 1a, in particular by the "explaining away" property of the unobserved $\mathbf{U}_t$. Marginalizing out $\mathbf{x}_t$ from the left-hand side of Eq. 9a and from Eq. 9f shows that that $p(\mathbf{y}_t|\mathbf{u}_{t-1}) = p(\mathbf{y}_t|\mathbf{y}_0, \ldots, \mathbf{y}_{t-1})$, and therefore $p(\mathbf{x}_t|\mathbf{u}_t) = p(\mathbf{x}_t|\mathbf{y}_0, \ldots, \mathbf{y}_t)$, which is Eq. 8. This completes the induction and therefore the proof of the theorem. ■

2.3 Training the rEFH

The proof shows that optimal inference—what is expressed by the information equality on the right side of Eq. 4—can be achieved, in density estimators that can be expressed as identifiable mixture models, by establishing the equality on the left side, $q(\mathbf{v}_t; \boldsymbol{\theta}) = p(\mathbf{v}_t; \boldsymbol{\theta})$, where $\mathbf{v}_t$ is constructed according to Eq. 3. To establish that equality, we train an exponential-family harmonium (EFH, depicted in Fig. 2a, and introduced by Welling et al., 2005) on these "augmented" data vectors. The joint distribution then takes the form of a Boltzmann distribution. For simplicity, we consider only exponential families whose sufficient statistics are the identity function. In accordance with Eq. 3, the visible layer consists of the current observations, $\mathbf{Y}_t$, and the previous hidden-layer sufficient statistics, $\mathbf{U}_{t-1}$; see Fig. 2b. Thus the joint distribution is

$$q(\mathbf{z}_t, \mathbf{u}_{t-1}, \mathbf{y}_t; \boldsymbol{\theta}) \propto h(\mathbf{u}_{t-1}) h(\mathbf{y}_t) h(\mathbf{z}_t) \exp\left\{\mathbf{z}_t{}'(W_{yz}\mathbf{y}_t + W_{uz}\mathbf{u}_{t-1} + \mathbf{b}_z) + \mathbf{b}_y'\mathbf{y}_t + \mathbf{b}_u'\mathbf{u}_{t-1}\right\}. \quad (10)$$

Here the functions h are the base measures, determined by the choices of exponential families. The intractable normalizer for q has been omitted. The corresponding conditional distributions are

$$q(\mathbf{z}_t|\mathbf{u}_{t-1}, \mathbf{y}_t; \boldsymbol{\theta}) = \frac{1}{\mathcal{Z}_z(\mathbf{y}_t, \mathbf{u}_{t-1})} h(\mathbf{z}_t) \exp\left\{(W_{yz}\mathbf{y}_t + W_{uz}\mathbf{u}_{t-1} + \mathbf{b}_z)'\mathbf{z}_t\right\}, \quad (11a)$$

$$q(\mathbf{y}_t|\mathbf{z}_t; \boldsymbol{\theta}) = \frac{1}{\mathcal{Z}_y(\mathbf{z}_t)} h(\mathbf{y}_t) \exp\left\{(W_{yz}'\mathbf{z}_t + \mathbf{b}_y)'\mathbf{y}_t\right\}, \quad (11b)$$

$$q(\mathbf{u}_{t-1}|\mathbf{z}_t; \boldsymbol{\theta}) = \frac{1}{\mathcal{Z}_u(\mathbf{z}_t)} h(\mathbf{u}_{t-1}) \exp\left\{(W_{uz}'\mathbf{z}_t + \mathbf{b}_u)'\mathbf{u}_{t-1}\right\}. \quad (11c)$$

The conditional distribution for the visible units has been separated into two conditionals, Eqs. 11b and 11c, corresponding to the observations and the previous hidden-layer means, respectively. From the graphical model, Fig. 2b, their product forms the conditional distribution over the visible units: $q(\mathbf{u}_t, \mathbf{y}_t|\mathbf{z}_t; \boldsymbol{\theta}) = q(\mathbf{u}_t|\mathbf{z}_t; \boldsymbol{\theta}) q(\mathbf{y}_t|\mathbf{z}_t; \boldsymbol{\theta})$.

The data distribution over $\mathbf{Y}_0, \ldots, \mathbf{Y}_T$ is assumed to be unknown, but the conditional distribution it induces over the hidden-layer sufficient statistics is known and deterministic:

$$p(\mathbf{u}_t|\mathbf{u}_{t-1}, \mathbf{y}_t; \boldsymbol{\theta}) := \delta(\mathbf{u}_t - \mathbb{E}_q[\mathbf{Z}_t|\mathbf{u}_{t-1}, \mathbf{y}_t; \boldsymbol{\theta}]). \quad (12)$$

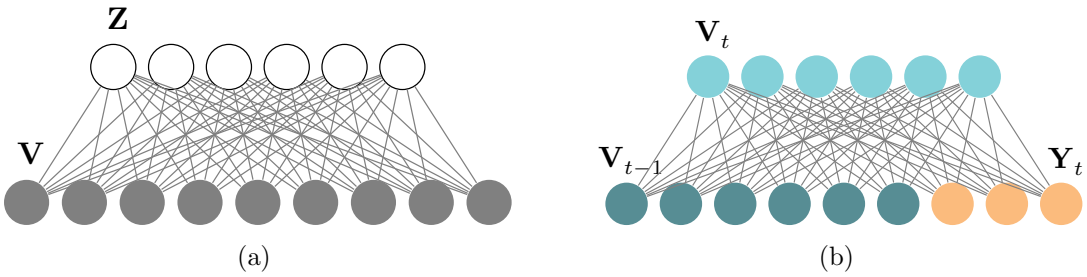

Figure 2: (a) A graphical model for a generic EFH. (b) The EFH composing time step t of the rEFH, i.e., one time slice in Fig. 1b. Each input vector consists of the sufficient statistics for the previous hidden-layer ($\mathbf{u}_{t-1} = \bar{\mathbf{z}}_{t-1}$, dark turquoise) and the current observations ($\mathbf{y}_t$, orange).

In words, the distribution over the current hidden-layer sufficient statistics is a delta function located at the posterior mean. In the EFH, this is computed as a deterministic "upward" (visible-to-hidden) pass through the weight matrix, followed by the pointwise nonlinearity appropriate for the exponential family to which the posterior distribution belongs (see Welling et al., 2005). We emphasize that, despite being parameterized (and by parameters of the model, no less), we treat the conditional distribution in Eq. 12 as part of the data distribution (p) because it provides part of the training data for the EFH.

Exact density estimation in an EFH—that is, gradient descent in the KL divergence between $p(\mathbf{v}_t;\boldsymbol{\theta})$ and $q(\mathbf{v}_t;\boldsymbol{\theta})$—requires prolonged Gibbs sampling, which introduces variance into the learning procedure as well as slowing it down (Hinton, 2002). We therefore proceed by stochastic gradient descent of the approximate "n-step contrastive divergence" (CD$_n$) objective function, which has $q = p$ at its minimum (Hinton, 2002; Hinton et al., 2006). At the cost of this approximation, contrastive-divergence learning buys fast, low-variance learning, since it requires only n (which can be small) full steps of Gibbs sampling from the model before the weights can be updated. We also ignore the dependence of the data distribution, $p(\mathbf{v}_t;\boldsymbol{\theta})$, on the parameters, introduced by Eq. 12, an approximation that we discuss further in Section 5, below.

The learning rules are therefore the standard rules for the EFH under one-step contrastive divergence (see Welling et al., 2005), applied to our data:

$$\Delta W_{yz} \propto \sum_{t=0}^{T} \left\langle \mathbf{Z}_t {\mathbf{Y}_t}' - \left\langle \hat{\mathbf{Z}}_t {\hat{\mathbf{Y}}_t}' \right\rangle_{q(\hat{\mathbf{z}}_t|\hat{\mathbf{u}}_{t-1},\hat{\mathbf{y}}_t)q(\hat{\mathbf{u}}_{t-1},\hat{\mathbf{y}}_t|\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)},$$

$$\Delta W_{uz} \propto \sum_{t=0}^{T} \left\langle \mathbf{Z}_t {\mathbf{U}_{t-1}}' - \left\langle \hat{\mathbf{Z}}_t {\hat{\mathbf{U}}_{t-1}}' \right\rangle_{q(\hat{\mathbf{z}}_t|\hat{\mathbf{u}}_{t-1},\hat{\mathbf{y}}_t)q(\hat{\mathbf{u}}_{t-1},\hat{\mathbf{y}}_t|\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)},$$

$$\Delta \mathbf{b}_z \propto \sum_{t=0}^{T} \left\langle \mathbf{Z}_t - \left\langle \hat{\mathbf{Z}}_t \right\rangle_{q(\hat{\mathbf{z}}_t|\hat{\mathbf{u}}_{t-1},\hat{\mathbf{y}}_t)q(\hat{\mathbf{u}}_{t-1},\hat{\mathbf{y}}_t|\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)},$$

$$\Delta \mathbf{b}_y \propto \sum_{t=0}^{T} \left\langle \mathbf{Y}_t - \left\langle \hat{\mathbf{Y}}_t \right\rangle_{q(\hat{\mathbf{y}}_t|\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)},$$

$$\Delta \mathbf{b}_u \propto \sum_{t=0}^{T} \left\langle \mathbf{U}_{t-1} - \left\langle \hat{\mathbf{U}}_{t-1} \right\rangle_{q(\hat{\mathbf{u}}_{t-1}|\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)}, \tag{13}$$

where the angle brackets indicate sample averages rather than true expectations.

3. Alternative extensions of the EFH to temporal data.

Sutskever and colleagues introduced an extension of the restricted Boltzmann machine (RBM) to temporal data, and proposed two different methods for training it (Sutskever and Hinton, 2007; Sutskever et al., 2009). We re-present them here to contrast them with the rEFH, and to show that the proof of the preceding section does not apply to these methods. For simplicity, we confine ourselves to the special case of the TRBM where the only direct dependence of an RBM on any of its predecessors in the chain is on the *hidden* units of its *immediate* predecessor (Fig. 3a). (We consider the more general model in the discussion below.) We derive both models in terms of an EFH, rather than an RBM as in the original papers, because the EFH is more general and can be accommodated within the framework of the cited works with minor adjustments only. Finally, our development of the models is slightly different, although the final result is the same.

Sutskever and Hinton (2007) define the following joint distribution:

$$q(\mathbf{z}_0, \mathbf{y}_0, \ldots, \mathbf{z}_T, \mathbf{y}_T; \boldsymbol{\theta}) = \prod_{t=0}^{T} q(\mathbf{z}_t, \mathbf{y}_t | \mathbf{z}_{t-1}; \boldsymbol{\theta}),$$

reflecting the independence properties of the graph in Fig. 3a. The graph also asserts the independence of $\mathbf{Y}_t$ and $\mathbf{Z}_{t-1}$, given $\mathbf{Z}_t$. This follows, straightforwardly, from the definition of the central component of the model,

$$q(\mathbf{z}_t, \mathbf{y}_t | \mathbf{z}_{t-1}; \boldsymbol{\theta}) := \frac{1}{\mathcal{Z}(\mathbf{z}_{t-1})} h(\mathbf{y}_t) h(\mathbf{z}_t) \exp\left\{ \mathbf{z}_t{}' W_{yz} \mathbf{y}_t + \mathbf{b}_y' \mathbf{y}_t + (W_{uz} \mathbf{z}_{t-1} + \mathbf{b}_z)' \mathbf{z}_t \right\},$$

since the components of the right-hand side that reference $\mathbf{z}_{t-1}$ and $\mathbf{y}_t$ can be factored into separate pieces. This equation has the form of a harmonium; that is, it can be produced by two conditional distributions of the form

$$\begin{aligned} q(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{y}_t) &= \frac{1}{\mathcal{Z}_z(\mathbf{y}_t, \mathbf{z}_{t-1})} h(\mathbf{z}_t) \exp\left\{ (W_{yz} \mathbf{y}_t + W_{uz} \mathbf{z}_{t-1} + \mathbf{b}_z)' \mathbf{z}_t \right\}, \\ q(\mathbf{y}_t | \mathbf{z}_t, \mathbf{z}_{t-1}) &= \frac{1}{\mathcal{Z}_y(\mathbf{z}_t)} h(\mathbf{y}_t) \exp\left\{ (W_{yz}' \mathbf{z}_t + \mathbf{b}_y)' \mathbf{y}_t \right\} = q(\mathbf{y}_t | \mathbf{z}_t). \end{aligned}$$

As it stands, learning in this model is intractable, since inference does not admit of a recursive procedure (like the Kalman filter). In order apply such a procedure, the authors decouple the EFHs through time with a gross approximation. They define a recursion $\mathbf{u}_t = \mathbb{E}_q[\mathbf{Z}_t | \mathbf{u}_{t-1}, \mathbf{y}_t]$, where the expectation is taken under the first conditional. Then they simply substitute $\mathbf{u}_{t-1}$ for $\mathbf{z}_{t-1}$. Thus, the approximate model is defined by the joint distribution

$$q(\mathbf{z}_t, \mathbf{y}_t | \mathbf{u}_{t-1}; \boldsymbol{\theta}) := \frac{1}{\mathcal{Z}(\mathbf{u}_{t-1})} h(\mathbf{y}_t) h(\mathbf{z}_t) \exp\left\{ \mathbf{z}_t{}' W_{yz} \mathbf{y}_t + \mathbf{b}_y' \mathbf{y}_t + (W_{uz} \mathbf{u}_{t-1} + \mathbf{b}_z)' \mathbf{z}_t \right\}, \tag{15}$$

corresponding to the graphical model in Fig. 3b; and the augmented data distribution is, as for the rEFH, defined by Eq. 12. This is the defining equation for the TRBM and the RTRBM.

Applying the substitution ($\mathbf{u}_{t-1}$ for $\mathbf{z}_{t-1}$) to the conditional distributions yields precisely Eq. 11a and Eq. 11b. This makes the deterministic recursion the same for both models, to wit, Eq. 12. Thus, the only difference between the (R)TRBM and the rEFH is that the latter additionally defines a conditional distribution over the previous hidden means, Eq. 11c. This can also be seen in the joint distributions: the EFH defined in Eq. 15 is *conditioned* on $\mathbf{u}_{t-1}$, rather than defining a (joint) distribution over it, as in Eq. 10. *The (R)TRBM does not define a distribution (q) over* $\mathbf{U}_t$, although the recursion extends the *data* distribution (p) to $\mathbf{U}_t$. The consequence can be seen in the subtle but important difference in the graphical models: the connections from $\mathbf{U}_{t-1}$ to $\mathbf{Z}_t$ are directed in the (R)TRBM (Fig. 3b), rather than undirected as in the rEFH (Fig. 1b). Essentially, the vector $\mathbf{U}_{t-1}$ is treated in the (R)TRBM like a (dynamic) bias rather than data, as in the rEFH; it is not allowed to "fantasize" about the past, only the present. This leads to different learning rules, which we exhibit below.

Treating the past hidden-unit activities as data is, however, required for our proof that the rEFH provides optimal inference (Section 2.2). Recall the basic idea: The hidden-unit activities of a(n identifiable) generative model become good summaries for the data they are required to generate; if the past hidden-unit activities are among those data, the summaries accumulate recursively. It is not obvious, therefore, that such a proof can be extended to the (R)TRBM (the left-hand equality in Eq. 4 will not be satisfied). Nevertheless, the RTRBM does appear to achieve nearly optimal inference, as we show in our experiments below. But choosing to treat the previous hidden-unit activities as data (the rEFH) does turn out to have a very important consequence: the dependence of these "data" on the rEFH's parameters, introduced by Eq. 12, can be safely ignored. When the recurrent activities are treated as a dynamic bias, on the other hand, ignoring their dependence on the parameters (the TRBM) produces suboptimal inference; while including the dependence in the learning rules (the RTRBM) introduces backpropagation-through-time. We show this in the following sections and argue why it is so in Section 5.

3.1 The TRBM

There are two training procedures for the graphical model of Fig. 3b. The first employs the approximation we have made use of for the rEFH. In minimizing KL divergence between data, $p(\mathbf{y}_t;\boldsymbol{\theta})$, and model, $q(\mathbf{y}_t|\mathbf{u}_{t-1};\boldsymbol{\theta})$, the dependence $\mathbf{u}_{t-1}$ on the parameters (Eq. 12) is ignored. Applying the method of one-step contrastive divergence to the distribution in Eq.

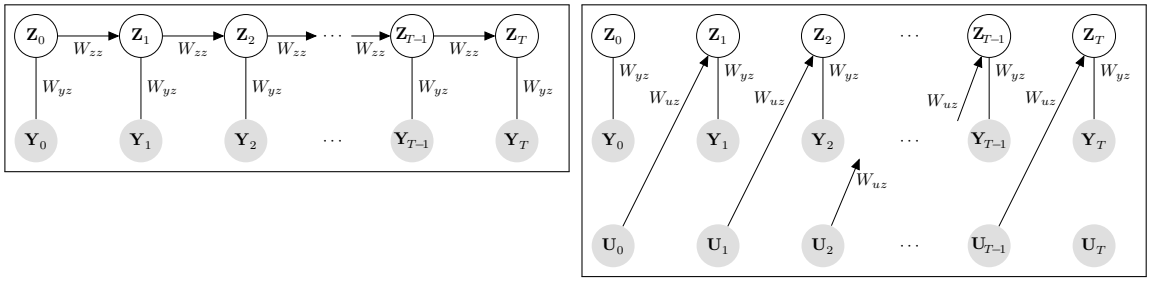

(a) (b)

Figure 3: Temporal extensions of the EFH. Each node represents a layer, within which there are no connections, but which has been collapsed to save space. (a) The non-approximated TRBM, to which learning rules are never directedly applied. (b) The TRBM and RTRBM. At each time step t, an EFH with dynamical bias $\mathbf{u}_{t-1} = \bar{\mathbf{z}}_{t-1}$ is trained on observations $\mathbf{y}_t$.

15 then yields

$$\begin{aligned}
\Delta W_{yz} &\propto \sum_{t=0}^{T} \left\langle \mathbf{Z}_t {\mathbf{Y}_t}' - \left\langle \hat{\mathbf{Z}}_t {\hat{\mathbf{Y}}_t}' \right\rangle_{q(\hat{\mathbf{z}}_t|\mathbf{u}_{t-1},\hat{\mathbf{y}}_t)q(\hat{\mathbf{y}}_t|\mathbf{u}_{t-1},\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)}, \\
\Delta W_{uz} &\propto \sum_{t=0}^{T} \left\langle \mathbf{Z}_t {\mathbf{U}_{t-1}}' - \left\langle \hat{\mathbf{Z}}_t {\mathbf{U}_{t-1}}' \right\rangle_{q(\hat{\mathbf{z}}_t|\mathbf{u}_{t-1},\hat{\mathbf{y}}_t)q(\hat{\mathbf{y}}_t|\mathbf{u}_{t-1},\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)}, \\
\Delta \mathbf{b}_z &\propto \sum_{t=0}^{T} \left\langle \mathbf{Z}_t - \left\langle \hat{\mathbf{Z}}_t \right\rangle_{q(\hat{\mathbf{z}}_t|\mathbf{u}_{t-1},\hat{\mathbf{y}}_t)q(\hat{\mathbf{y}}_t|\mathbf{u}_{t-1},\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)}, \\
\Delta \mathbf{b}_y &\propto \sum_{t=0}^{T} \left\langle \mathbf{Y}_t - \left\langle \hat{\mathbf{Y}}_t \right\rangle_{q(\hat{\mathbf{y}}_t|\mathbf{u}_{t-1},\mathbf{z}_t)} \right\rangle_{q(\mathbf{z}_t|\mathbf{u}_{t-1},\mathbf{y}_t)p(\mathbf{u}_{t-1},\mathbf{y}_t)}.
\end{aligned} \tag{16}$$

We emphasize the subtle difference between Eq. 13 and Eq. 16; in particular, no "updated" version of the posterior means, $\hat{\mathbf{u}}_{t-1}$, ever appears in the latter.

3.2 The RTRBM

If the dependence of the KL divergence on the parameters via the data distribution, introduced by Eq. 12, is *not* ignored (Sutskever et al., 2009), the learning rules accrue an extra term, which must be computed with backpropagation-through-time (BPTT). Thus the RTRBM learning rules are

$$\Delta W_{yz} = (\Delta W_{yz})_{\text{TRBM}} + \eta \sum_{t=0}^{T} \mathbf{g}_t {\mathbf{y}_t}', \qquad \Delta W_{uz} = (\Delta W_{uz})_{\text{TRBM}} + \eta \sum_{t=0}^{T} \mathbf{g}_t {\mathbf{u}_{t-1}}'$$

$$\Delta \mathbf{b}_z = (\Delta \mathbf{b}_z)_{\text{TRBM}} + \eta \sum_{t=0}^{T} \mathbf{g}_t, \qquad \Delta \mathbf{b}_y = (\Delta \mathbf{b}_y)_{\text{TRBM}};$$

$$\mathbf{g}_t = J_t' W_{uz}' \left(\frac{\partial \log q(\mathbf{y}_{t+1}|\mathbf{u}_t;\boldsymbol{\theta})}{\partial \mathbf{b}_z} + \mathbf{g}_{t+1} \right).$$

where the last line defines the backwards recursion for BPTT (for a derivation, see **Appendix B**), η is the learning rate, and J_t is the Jacobian of the nonlinearity (e.g., logistic function for Bernoulli units).[2]

4. Experiments

We compare performance of the rEFH, TRBM, and RTRBM on two different kinds of data: a linear dynamical system (LDS) that is observed via a "probabilistic population code" (PPC) (Ma et al., 2006); and the bouncing-ball data used by Sutskever and Hinton (2007) and Sutskever et al. (2009). For the LDS, we have reported previously a subset of the results for the rEFH (Makin et al., 2015). Here we compare the filtering performance of the rEFH with that of the TRBM and RTRBM on these data, and investigate the failure modes of the latter, especially in connection with the order of the dynamics that can be learned. For the bouncing-ball data, Sutskever (2013) characterized the performance of the TRBM and RTRBM in terms of prediction of the next frame. Here, we characterize the rEFH's performance in the same way on the same data.

4.1 Details of the training procedures

With the exception of the rEFHs trained on LDSs (Section 4.2), all models and experiments used the same training scheme, which was matched to that reported by Sutskever (2013). Weight changes were made according to 25-step contrastive-divergence with a fixed momentum of 0.9 and a learning rate declining linearly from 1/1500 (LDS) or 1/100 (bouncing balls)[3], on "minibatches" consisting of 100 time steps of a single trajectory. New batches of 400 trajectories (i.e., 400*100 = 40000 sample vectors) were created every 5 epochs, for 250 epochs. These models were initialized with 30 batches of static pre-training (i.e., setting recurrent activities to zero) at CD-5.

The rEFHs of Section 4.2, on the other hand, could be trained much more quickly. No static pre-training was required, only CD-1 was used throughout, and training terminated after only 90 epochs. Minibatches consisted of 40 vectors, each corresponding to the same time step of 40 different trajectories.[4] Batches, renewed every 5 epochs, consisted of 1000 time steps (i.e., 40*1000 = 40000 sample vectors). Learning rates declined *exponentially*, according to $1.1^{-\text{epoch}}$, from 1/500 (Poisson-Bernoulli weights), 1/50 (Bernoulli-Bernoulli weights) or 1/120 (biases). Momentum started at about 0.5 and exponentially approached 0.98 ($\rho = 0.98 - 0.5 * 1.1^{-\text{epoch}}$). A small weight decay of 0.001 was also used.

2. Of course, the nonlinearity acts element-wise, so the Jacobian is diagonal, and in practice is implemented as a Hadamard product rather than a matrix multiplication.
3. The differing learning rates reflect the size of the inputs: the LDS data are Poisson distributed with maximum mean firing rates of about 10, whereas the bouncing-ball "pixels" are Bernoulli distributed.
4. This minibatch scheme is precluded for the RTRBM by the use of BPTT, which requires many steps of the *same* trajectory.

4.2 Linear dynamical system with PPC observations

We aimed to quantify model performance precisely, so we constructed a generative model for the data (consistent with Fig. 1a) for which the posterior distribution can be computed in closed form—despite nonlinear, non-Gaussian emissions. The "backbone" of the model is a second-order linear dynamical system; that is, position and velocity depend on previous position and velocity. The observations at each time step are a product over fifteen conditionally independent Poisson "neurons" with fixed-width Gaussian-shaped tuning to the position—a "probabilistic population code" (Ma et al., 2006). This could be thought of as a crude model of sensory responses to a moving object. Since the tuning curves evenly tile the space, the posterior distribution over the hidden dynamical state is in fact (approximately) a normal distribution (Ma et al., 2006). This makes the filtering distribution (the RHS of Eq. 5) computable in closed form—the Kalman filter (KF)—where, in place of the observation and the emission variance, respectively, one uses the center of mass of the population code and the tuning-curve width scaled by the total number of spikes (Ma et al., 2006). This model can be learned with an EM algorithm (Ghahramani and Hinton, 1996), which allows us to benchmark the rEFH, TRBM, and RTRBM against easily interpreted quantities. In particular, we can consider Kalman filters applied to models learned by EM under the assumption of first- ($\mathrm{KF}^1_{\mathrm{EM}}$) or second-order ($\mathrm{KF}^2_{\mathrm{EM}}$) dynamics. We also considered a zeroth-order model (KF^0), which assumes no dynamics at all; and a KF that has access to the parameters of the true dynamical system ($\mathrm{KF}^2_{\mathrm{OPT}}$—no learning required). We chose these parameters so that any first-order approximation will be poor, allowing us to distinguish these cases from each other.

Fresh data were generated for testing: for each model the posterior mean, $\mathbb{E}_q[\mathbf{Z}_t|\mathbf{v}_t]$, was computed at each time step. From each vector of hidden activities, the position of the underlying "stimulus" was decoded (using the technique we have introduced for similar, static models; see Makin et al. 2013 or **Appendix C**) and used to compute an error with respect to the actual position. We did the same using the posterior mean from Kalman filters learned using $\mathrm{KF}^1_{\mathrm{EM}}$, $\mathrm{KF}^2_{\mathrm{EM}}$; the zeroth-order model; and the "optimal" (no-learning) model. For each, we report the mean, across time steps (1000) and trajectories (40), of the square of these errors (MSE). In order to characterize susceptibility of training to local minima, 20 of each of the EM-based models ($\mathrm{KF}^1_{\mathrm{EM}}$, $\mathrm{KF}^2_{\mathrm{EM}}$) were trained from scratch. Likewise, 20 of each of the rEFH, TRBM, and RTRBM were trained from scratch for each of 20 different hidden-layer sizes. The distribution of the resulting MSEs are shown in the box-and-whisker plot in Fig. 4. Also indicated with lines are the performances of the *best* of each of the $\mathrm{KF}^1_{\mathrm{EM}}$ and $\mathrm{KF}^2_{\mathrm{EM}}$ models, and the performances of the models requiring no training, KF^0 and $\mathrm{KF}^2_{\mathrm{OPT}}$. (Since a single set of testing data was used for all models, there is no variation in performance—or, therefore, box plots—for the no-training models.)

Evidently, the TRBM learns only first-order dynamics, as suggested by Sutskever (2013), who reported that it tends to generate data that move in a random walk. The best RTRBMs approach the performance of the best second-order Kalman filter ($\mathrm{KF}^2_{\mathrm{EM}}$, the order of the true underlying dynamics). The rEFHs are on average slightly worse, but given at least 195 hidden units, the best are very close to the best RTRBMs, which do not appear to improve after about 165 hidden units. The rEFHs are certainly able to learn second-order dynamics, unlike the TRBM; and they do not require BPTT, unlike the RTRBM. In fine,

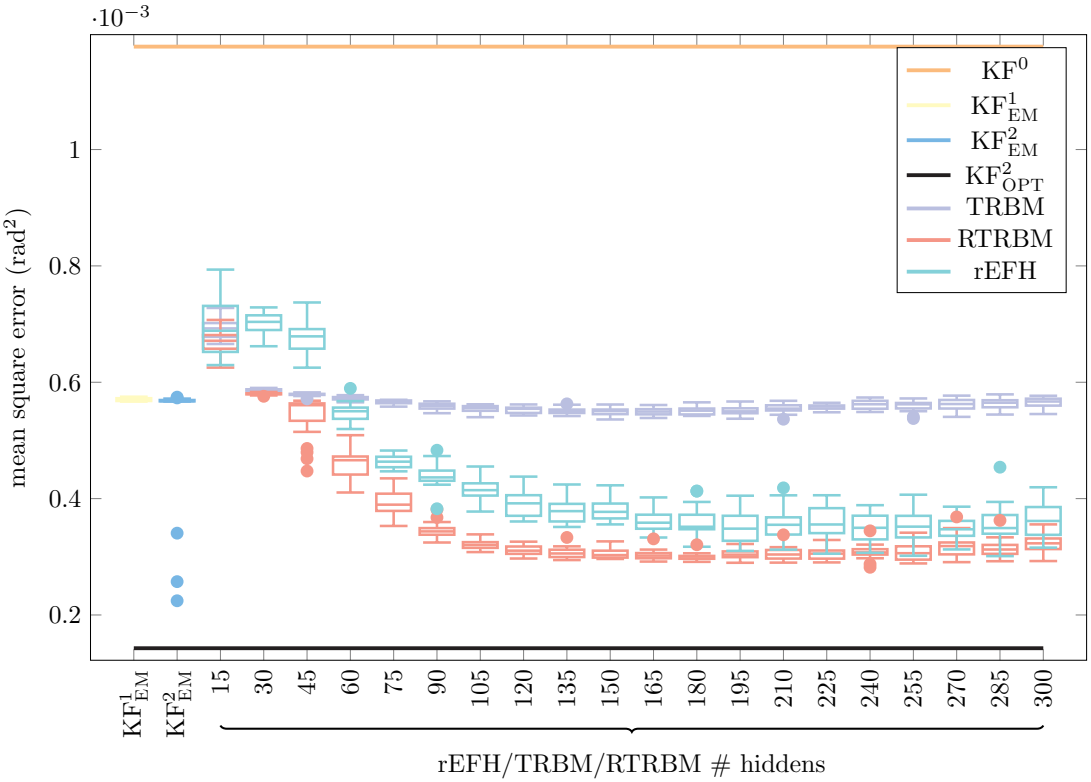

Figure 4: Box-and-whisker plots of mean square errors (MSEs) for the rEFH, TRBM, RTRBM, and benchmark filters, across 20 repeated training instances, on a single testing set. Median MSE for each set of networks is marked with a line; the box contains the interquartile range; whiskers extend to 1.5× the interquartile range, beyond which outliers are marked with filled circles. Performance of the benchmark filters KF^{0}and KF$^2_{\text{OPT}}$ is marked with long horizontal lines: no boxplot is possible, because these models require no training. We have reported previously similar results for the rEFH and the benchmarks (Makin et al., 2015).

the performance deficit of the rEFH relative to the RTRBM appears not to be insuperable, as with the TRBM, but rather to be remediable—at least on these data—with an increase in hidden units. This cost may be offset to some extent by the savings accrued from not needing BPTT; we explore this below on a more difficult task, the bouncing-ball data set. Finally, we note that only outliers among the KF$^2_{\text{EM}}$ models actually achieve performance superior to that of the median rEFH or RTRBM, as long as these neural networks have at least 60 hidden units.

4.3 Bouncing balls

Each data vector consists of a 30×30 image patch containing three "balls" (circles) that bounce off each other and the walls with complete energy conservation. The nonlinear

dynamics greatly complicate computation of the posterior distribution over the hidden state—the positions and velocities of the balls—so we compute from the hidden units not the optimal posterior means, but the best predictions of the next frame, as proposed by Sutskever (2013). We note first of all that the *current frame* predicts the next frame with a per-pixel, mean square error (MSE) of 0.015.

Our implementation of the TRBM, with 400 hidden units, achieves a MSE of 0.046, close to the 0.04 reported by Sutskever (2013). Interestingly, this is higher (worse) than the current-frame prediction, which suggests that this TRBM does not learn any dynamical model at all. In fact, this is consistent with the results above: The TRBM can learn only first-order dynamics; and for the bouncing balls, the optimal first-order model is the zeroth-order model, since the velocities are indeed constant (until collisions), and the *average*, over all trajectories, of the next position is the current position (velocities are on average zero). That the TRBM fails even to learn even the zeroth-order model (the current-frame predictions) suggests, perhaps, a limitation imposed by the nonlinearities (collisions) in the dynamics.

Our implementation of the RTRBM, with 400 hidden units, achieves a MSE of 0.008, again close to the results reported by Sutskever (2013), in this case 0.007. The rEFH with 400 hidden units achieves a per-pixel MSE of 0.014: superior to the TRBM, as well as the zeroth- and first-order predictions, but inferior to the RTRBM. However, training time is also much shorter (see Table 1), since BPTT is not required.[5] Increasing the number of hidden units to 625 increases the rEFH's training time to that of the RTRBM, about 2.5 hours, and brings its MSE closer, down to 0.010 MSE. Allowing a little under five hours accommodates 1000 hidden units in the rEFH and approximately matches the RTRBM's MSE at 0.008.

model	# hidden	MSE	training time (hrs)
0^{th}-order	—	0.015	—
TRBM	400	0.046	1.43
TRBM	625	0.040	2.53
RTRBM	400	0.008	2.56
rEFH	400	0.014	1.42
rEFH	625	0.010	2.56
rEFH	1000	0.008	4.76

Table 1: For the bouncing-ball data, a comparison of next-frame per-pixel MSE, and training times.

5. The comparison is between our own implementations of the RTRBM and the rEFH. Our highly vectorized, GPU-optimized MATLAB implementations take on the order of hours; on the same machine, Sutskever's original PYTHON implementation of the RTRBM takes a few days. All of our code is publicly available at `https://github.com/jgmakin`.

4.4 Generating trajectories

The rEFH enjoys a different advantage over the (R)TRBM: generating sequences is orders of magnitude less costly. Because the rEFH defines a distribution over the previous hidden units, generation in the reverse direction is particularly simple: from each hidden vector, a "previous" hidden vector as well as the current observation can be drawn immediately, and the process iterated. In the (R)TRBM, by contrast, generating sequences (in either direction) requires at each time step repeated Gibbs sampling of the hidden and visible units, conditioned on the previous hidden means, before the draws converge to the stationary distribution. In practice, this appears to be about 50 steps—i.e., 50 upward and 50 downward passes (Sutskever, 2013). Examples of bouncing-ball trajectories generated by the rEFH can be found in the supplemental material.

5. Conclusions

The "recurrent exponential-family harmonium" (rEFH) generalizes EFHs to time-series data. Its central motivation is to provide optimal inference to the hidden state of "the world" that produces those data, even when the world's true generative model is intractable to Bayesian inversion, hard to express in closed-form equations, or even unknown. Thus, in contrast to directed graphs like dynamic Bayes Nets (Murphy, 2002), where one attempts to formulate the true generative model and then derive an inference algorithm for it, we relaxed the assumption that the generative model of the world and the rEFH correspond, in favor of a weaker one about identifiability: The rEFH's and the world's marginal distributions over observations must be expressible in terms of the same identifiable finite mixture model (see **Appendix A** for technical details). Consequently, the sufficient statistics for the true hidden state are not derived in closed form (like, e.g., the posterior cumulants recursively computed in a Kalman filter), but are nevertheless guaranteed to be encoded in the hidden units of the trained rEFH. The idea behind the proof of this guarantee is simple: good density estimators make their hidden units—or rather, the sufficient statistics for them—good summaries of their input data. When such summaries are also included among the inputs—at the next time step (Fig. 1)—the succeeding hidden vector is forced to summarize both the new inputs and the old summaries. This recursive procedure builds up a master summary that takes into account all relevant past information.

The other conspicuous consideration regarding the choice of directed and undirected (or at least EFH-like) models is the trade-off between inference and generation. In directed models, generation is easy, but inference—in all but a few special cases—is hard, and approximate techniques are necessary; whereas in the EFH-like undirected models, inference (to its own hidden states) is easy, but generation is hard, because it requires Gibbs sampling. But in the rEFH, generating sample trajectories—in reverse time order—is easy, requiring only a single pass through the model. This is not the case in the alternative architecture proposed by Sutskever (2013), the TRBM and RTRBM (Fig. 3b), where prolonged Gibbs sampling is required.

In practice, some of the assumptions of the proof are mildly violated: e.g., that the model has become a perfect model for the training data. To show that the proof is robust against such violations, we tested the model experimentally on data from a generative model that has been contrived to yield closed-form solutions to the inference problem, a

second-order system whose dynamics are, and whose emissions can be re-formulated to be, linear-Gaussian. This enabled us to demonstrate conclusively that the rEFH learns second-order dynamics from reports of position only—even though its recurrent connections span only one time step.

But the most interesting approximation is the use of a sub-optimal learning rule. During learning, changes to the model's parameters ramify into the data, since they consist in part of the model's own hidden-unit means (Eq. 12). We ignored this, which amounts to assuming that whatever improvements to the model are effected by a parameter update are not (entirely) undone by the ensuing changes to the data. This obviates the need for backprop through time (BPTT). In Sutskever's architecture (Fig. 3b), the same approximation turns the RTRBM into a TRBM, and loses the ability to learn dynamics beyond first-order, at least without increasing the span of its recurrent connections. In contrast, the only advantage the RTRBM appears to enjoy over the rEFH is more economical use of hidden units (Fig. 4, Table 1), which comes at the price of slower training times (BTPP being required), and the inability to generate trajectories in one pass. Some of the rEFH's training-time advantage can be bartered for performance by increasing the size of its hidden layer —although to eliminate completely the performance gap on Sutskever's bouncing-ball data set ultimately requires about a factor of 2 *more* time for the rEFH (Table 1). But for data with trajectories longer than 100 time steps, the advantage of neglecting BPTT is even greater, and the relative merits of the rEFH increase.

Why can BPTT be neglected for the rEFH but not the (R)TRBM? The answer does not follow obviously from application of BPTT to the rEFH architecture, which yields non-zero terms. We consider this the most important theoretical question for future work. Second, empirically, the accumulation of master summaries described above is ultimately limited, obviously, by the capacity of the hidden layer, so not all past history can be "summarized" in the way just described. (This is consistent with the proof because the capacity limits would likewise limit the generative fidelity of the model to the data, compromising one of the proof's requirements.) Finally, it also seems unlikely that the model, without modification, could learn "pathologically" long dependencies. Thus, although the toy data sets we picked allowed us to benchmark the models, it remains to determine how the model performs on more challenging time-series data.

Acknowledgments

Funding was provided by DARPA (N66001-10-C-2010, W911NF-14-2-0043), NIH Vision Training Grant (5T32EY007120-22), the UCSF Wheeler Center for the Neurobiology of Addiction, and (BKD) the National Science Foundation Graduate Research Fellowship under Grant No. 1144247. Some of the EFHs were trained using Tesla K40 GPUs, the generous donation of the Nvidia Corporation.

Appendix A. Information retention and identifiability

The essential result used to prove optimal inference in the rEFH comes from Makin and Sabes (2015), who showed that for a certain class of latent-variable density estimators,

depicted in Fig. 5, the following implication holds:

$$q(\mathbf{v}_t) = p(\mathbf{v}_t) \implies p(\{\mathbf{x}\}_0^t|\mathbf{v}_t) = p(\{\mathbf{x}\}_0^t|\mathbf{u}_{\mathbf{Z}_t}(\mathbf{v}_t)) \implies p(\mathbf{x}_t|\mathbf{v}_t) = p(\mathbf{x}_t|\mathbf{u}_{\mathbf{Z}_t}(\mathbf{v}_t)). \quad (17)$$

Time-dependence is not discussed in that paper, but we have subscripted the variables in Eq. 17 with t to facilitate our application of that result to the models of the present paper. The equation can be read as saying that a match between the model and data marginals (the left-hand equation) implies that the posterior distribution over all true hidden states up to time t, $\{\mathbf{x}\}_0^t$, is the same whether conditioning on the observed data at that time step, $\mathbf{v}_t$, or on any set of sufficient statistics, $\mathbf{u}_{\mathbf{Z}_t}(\mathbf{v}_t)$, for the *model*'s hidden states, $\mathbf{Z}_t$. For the EFH, the sufficient statistics consist of the vector of hidden-unit means. (We have used "U" rather than the usual "T" for sufficient statistics because the latter is used for final time.) The final equation says that the same applies (*a fortiori*) to the posterior distribution over just the *current* true hidden state, $\mathbf{X}_t$. More colloquially, we can say that training the model to be a good density estimator guarantees that no information about the true hidden state will be lost when making a (deterministic) "upward" pass through the model, from the visible to hidden units.

In the experiments, and in the proof in the main text, the data vector $\mathbf{v}_t$ consists of the current observations, concatenated with the previous hidden-unit means. Our exposition of identifiability is greatly facilitated, however, by working with *samples*, rather than means, from those hidden units: The EFH, after all, produces Bernoulli samples, not real numbers, on its "left side" (see Fig. 2b). We justify this approximation by noting that, for sufficiently long hidden vectors, little information will be lost even from a *sample* from the vector of hidden means. Indeed, when sample rather than mean vectors were used in experiments like those in the main text, there was no noticeable performance difference (Makin et al., 2015). Thus we replace Eq. 3 in the main text with

$$\mathbf{v}_t := \begin{bmatrix} \mathbf{z}_{t-1} & \mathbf{y}_t \end{bmatrix}. \quad (18)$$

Previously (Makin and Sabes, 2015), we provided sufficient conditions for membership in the class of models for which Eq. 17 holds. Here we consider those technical conditions with respect to the dynamical models discussed in the present paper. We discuss both the specific model on which the experiments (especially those of Section 4.2) were performed, and dynamical models more generally. Although we argue below that satisfaction of the conditions is plausible, we suspect that in practice some of them are mildly violated; we indicate where below. The experiments show, what the proof does not, that the guarantees are robust to such violations, in the sense that mild departures from the assumptions lead only to mild departures from the conclusion, optimal inference.

Condition 1. The first technical condition is that the observations $\mathbf{V}_t$ be the same type of random variable in both the "world" (p, the "data distribution," in Hinton's terminology) and the "model" (q, "model distribution"). More precisely, the conditional distributions $p(\mathbf{v}_t|\{\mathbf{x}\}_0^t)$ and $q(\mathbf{v}_t|\mathbf{z}_t)$ (see again Fig. 5) must both be expressible in terms of some common conditional distribution $f(\mathbf{v}_t|\boldsymbol{\lambda}_t)$, with $\boldsymbol{\lambda}_t = \phi(\{\mathbf{x}\}_0^t)$ in the world, and $\boldsymbol{\lambda}_t = \psi(\mathbf{z}_t)$ in the model, for some functions $\phi(\cdot)$ and $\psi(\cdot)$. For example, if the observations were truly (p) generated by products over Poisson distributions, the model (q) observations must be expressible as products over Poissons, as well.

Figure 5: Probabilistic graphical models. Observed nodes are shaded. Training the model (q) makes the marginals match: $p(\mathbf{v}_t) = q(\mathbf{v}_t)$. (a) The "world's" generative process. (b) The model's generative process.

In real and artificial neural networks, this condition is not particularly restrictive: it says that neurons must respond with the same *kind* of statistics—Poisson, Bernoulli, Gaussian, etc.—whether driven bottom-up or top-down. In the experiments in the main text, the hidden units were Bernoulli random variables while the "sensory" units were Poisson. Using samples (rather than means) for the recurrent units, then, makes the function f a product over Bernoulli distributions and Poisson distributions. What about $\psi(\mathbf{z}_t)$ and $\phi(\{\mathbf{x}\}_0^t)$? For the model, $\psi(\mathbf{z}_t)$ was a "downward" (hidden-to-visible) pass through the weight matrix, followed by the element-wise nonlinearities: the logistic function for Bernoulli units, the exponential function for the Poisson units. For the "world," $\boldsymbol{\lambda}_t = \phi(\{\mathbf{x}\}_0^t)$ was more complicated. The "right-hand" subvector (cf. Eq. 18), for the Poisson units, consisted of the embedding of the current hidden state $\mathbf{x}_t$ into (Gaussian) tuning curves. The "left-hand" subvector, for the Bernoulli units, was created with the recursion defined by Eq. 12 (see also Fig. 1 in the main text). For simplicity in the following discussion, we substitute samples $\mathbf{z}_{t-1}$ for the mean vector $\mathbb{E}_q[\mathbf{Z}_{t-1}|\mathbf{v}_{t-1}; \boldsymbol{\theta}]$.

Condition 2. The second condition is the antecedent proposition in Eq. 17, sc., that the model density match the data density. Training the EFH as a density estimator—that is, changing its parameter so as to descend the gradient of a function with $q(\mathbf{v}_t) = p(\mathbf{v}_t)$ at its minimum—is supposed to guarantee this. It must be noted, however, that the proof in our previous work (Makin and Sabes, 2015) requires strict equality, whereas the use of n-step contrastive divergence with small n (which is believed to produce suboptimal density estimators; see Hinton, 2010) and the finite number of training samples suggests that the equality need only be approximate in practice.

Condition 3. The third condition is that the number of hidden-variable states be finite. Although the data in the experiments were in fact generated with real-valued hidden states, $\mathbf{X}_t$, we rely on the fact that only a finite number of such states was ever generated. This condition is in fact presupposed by the antecedent proposition, $q(\mathbf{v}_t) = p(\mathbf{v}_t)$: since the EFH has a finite, albeit enormous ($2^{\text{number of hidden units}}$), number of hidden states, the model distribution cannot be guaranteed to match the data distribution unless the latter likewise has a finite (albeit possibly enormous) number of hidden states.

Condition 4. The final condition is the most complicated. To begin with, we re-write the data and model distributions in terms of $\boldsymbol{\Lambda}_t$: see Fig. 6. Since the number of hidden states is finite, these can be thought of as finite mixture models. The required condition is that the mixture of the conditional distributions, $f(\mathbf{v}_t|\boldsymbol{\lambda}_t)$, be "identifiable," which is to

Figure 6: The models from Fig. 5 are rewritten in terms of the common conditional distribution $f(\mathbf{v}_t|\boldsymbol{\lambda}_t)$ and the prior distributions $p(\boldsymbol{\lambda}_t)$ and $q(\boldsymbol{\lambda}_t)$ that are induced by the changes of variables $\mathbf{\Lambda}_t := \phi(\{\mathbf{X}\}_0^t)$ and $\mathbf{\Lambda}_t := \psi(\mathbf{Z}_t)$, respectively. (a) The "world's" generative process. (b) The model's generative process.

say that the map from the parameters of the mixture to the marginal $p(\mathbf{v}_t)$ be injective (up to relabeling of the hidden units). In particular, in identifiable mixture models, $p(\mathbf{v}_t) = q(\mathbf{v}_t) \implies p(\boldsymbol{\lambda}_t) = q(\boldsymbol{\lambda}_t)$, which is the key to proving Eq. 17.

Identifiability in this sense is the rule rather than the exception (Titterington et al., 1985). Nevertheless, several exceptions are notable, among which is the RBM, a mixture of products of Bernoulli distributions (Titterington et al., 1985). On the other hand, a recent result (Allman et al., 2009) restores identifiability to many of these exceptions by relaxing the definition very mildly. In particular, Allman et al. (2009) consider "generically" identifiabe mixture models, that is to say, those that are unidentifiable only on sets of measure zero. Their results imply, among other things, that the RBM is generically identifiable as long as there are more than twice as many visible as hidden units. For the bouncing-ball data, then, the RBM can have up to 899 hidden units—each image being itself 900 pixels, yielding a total of 1799 visible (image plus recurrent) units.

For the EFH in our linear-dynamical-system experiments, application of the results of Allman is slightly more involved, since the conditional distribution $f(\mathbf{v}_t|\boldsymbol{\lambda}_t)$ is a product over both Bernoullis and Poissons. The relevant result is Theorem 4 in that paper, which supplies identifiability conditions for finite mixtures of finite-measure products. Consider in particular an EFH with: N hidden units, each with cardinality (number of possible states) n; M non-recurrent units, each with cardinality m; and therefore $N + M$ visible units, with the j^{th} hidden unit having cardinality k_j. Distributions defined by this model are generically identifiable if (Allman et al., 2009)

$$\min(r, \kappa_1) + \min(r, \kappa_2) + \min(r, \kappa_3) \geq 2r + 2, \tag{19}$$

where we have defined $r = n^N$, the number of hidden *states*; integers $\kappa_i := \prod_{j \in S_i} k_j$; and sets S_1, S_2, S_3, which constitute a disjoint, exhaustive partition of the visible units, $\mathbf{v}_t$.

Assuming $r > 2$, the optimal strategy for partitioning the visible units is to try to have $\kappa_1 = \kappa_2 = r$, allowing κ_3 to be whatever remains, since $\min(r, \kappa_3)$ is guaranteed to be at least 2 in any case. Thus, S_3 should contain just one visible unit, and S_1 and S_2 should each contain half the remainder, divided so that κ_2 and κ_2 are (roughly) equal. Under this scheme, we have $\kappa_2 = \kappa_3 \approx n^{N/2} m^{M/2}$. (We assume that N and M are much greater than 1, allowing us to ignore the one unit assigned to S_3.) Requiring that this be greater than $r = n^N$ is tantamount to requiring that

$$n^N < m^M. \tag{20}$$

What about our Poisson-Bernoulli EFH? Technically, the theorem applies only to finite-measure products, whereas the Poisson units can take on an infinite number of states. On the other hand, the conditional distribution of each unit can be approximated arbitrarily well by a sufficiently long, but truncated, version of the Poisson distribution. Under such an approximation, Eq. 19 is satisfied, each EFH is generically identifiable, and the perfectly trained rEFH is guaranteed (with probability 1) to provide optimal inference about the true hidden causes of the linear dynamical system—i.e., Eq. 17 holds.

On the other hand, identification of such a model would require an impractically large number of training samples. For any finite number of samples, almost all of the infinite number of states that a Poisson random variable can take on will have never be observed; thus, their probability will be indistinguishable from zero. In our experiments, we tested rEFHs with $M = 15$ Poisson units, and N between 15 and 300 hidden/recurrent Bernoulli units, having cardinality 2. If we choose κ_2 as above, satisfying Eq. 20 for, say, $N = 150$, requires that $m = 1000$, whereas the largest Poisson mean was about 10, making a nonzero count (in the number of samples generated for the experiments) essentially zero for any integer greater than about 30.

We have shown, then, that the model used in the linear-dynamical-system experiments is formally generically identifiable, in the sense of Allman et al. (2009). Nevertheless, in practice, the limited number of training samples will limit this identifiability, due to both condition 2 and condition 4. The fact that this does not prevent the trained model from providing good inference to the true hidden state suggests either that the training procedure is robust against this type of failure—that mild departures from identifiability result in mild departures from optimal inference—or that identifiability is not in fact required for Eq. 17 in this model. Indeed, it is not in general (Makin and Sabes, 2015), although no other set of sufficient conditions has been adduced. In either case, Eq. 20 suggests how to modify rEFHs in order to guarantee identifiability. Generally speaking, the non-recurrent units must be either more numerous, or able to take on more states, than the recurrent/hidden units.

Appendix B. Derivation of the update rules for the RTRBM

Since our update equations differ slightly from those given by Sutskever et al. (2009), and no explicit derivation is given in that work, we derive the training algorithm for the RTRBM here. (We use transposes—denoted $'$, since "T" is used for final time—to indicate the shape of our matrix derivatives. For now, $\boldsymbol{\theta}$ is a generic vector of parameters.) First, we consider a generic, deterministic, recurrent neural network (RNN) with inputs $\{\mathbf{Y}\}_0^T \sim p(\{\mathbf{y}\}_0^T)$ and "hidden"/output-unit activities $\bar{\mathbf{z}}_t$ recursively computed as

$$\bar{\mathbf{z}}_t = f(\bar{\mathbf{z}}_{t-1}, \mathbf{y}_t, \boldsymbol{\theta}). \tag{21}$$

We leave the loss function undetermined for now, but let it depend on the output at every time step, as well as directly on the parameters:

$$L(\bar{\mathbf{z}}_0(\boldsymbol{\theta}), \ldots, \bar{\mathbf{z}}_{T-1}(\boldsymbol{\theta}), \boldsymbol{\theta}). \tag{22}$$

To compute parameter changes, we take the (total) derivative of the average (under the data distribution) loss function:

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'}\langle L\rangle_{p(\{\mathbf{y}\}_0^T)} &= \left\langle \frac{\partial L}{\partial \boldsymbol{\theta}'} + \sum_{t=0}^{T} \frac{\partial L}{\partial \bar{\mathbf{z}}_t{}'}\frac{\mathrm{d}\bar{\mathbf{z}}_t}{\mathrm{d}\boldsymbol{\theta}'}\right\rangle_{p(\{\mathbf{y}\}_0^T)} \\ &= \left\langle \frac{\partial L}{\partial \boldsymbol{\theta}'} + \sum_{t=0}^{T} \frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}\frac{\partial\bar{\mathbf{z}}_t}{\partial\boldsymbol{\theta}'}\right\rangle_{p(\{\mathbf{y}\}_0^T)}.\end{aligned} \tag{23}$$

The first equality is just standard calculus; the second follows because $\boldsymbol{\theta}$ only occurs in the final terms in the chains of derivatives. Thus, one can either compute the total effect of $\boldsymbol{\theta}$ on each $\bar{\mathbf{z}}_t$, and then compute the direct effects of the latter on L; or one can compute the *direct* effect of $\boldsymbol{\theta}$ on each $\bar{\mathbf{z}}_t$, and then compute the *total* effect of the latter on L.

Because of Eq. 21, the effect on L of $\bar{\mathbf{z}}_t$, the current hidden state, is the sum of its direct effect and the effects caused by its influence on *future* hidden states, $\bar{\mathbf{z}}_{t+k}, k > 0$. This can be expressed recursively as

$$\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'} = \frac{\partial L}{\partial \bar{\mathbf{z}}_t{}'} + \frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_{t+1}{}'}\frac{\partial\bar{\mathbf{z}}_{t+1}}{\partial\bar{\mathbf{z}}_t{}'}. \tag{24}$$

The recursion is initialized at $\mathrm{d}L/\,\mathrm{d}\bar{\mathbf{z}}_T{}' = \partial L/\partial\bar{\mathbf{z}}_T{}'$ and run backwards in time. The backprop-through-time (BPTT) algorithm consists of Eqs. 23 and 24. Applying it to any particular model requires only working out four partial derivatives, $\partial L/\partial\bar{\mathbf{z}}_t{}'$, $\partial L/\partial\boldsymbol{\theta}'$, $\partial\bar{\mathbf{z}}_{t+1}/\partial\bar{\mathbf{z}}_t{}'$, and $\partial\bar{\mathbf{z}}_t/\partial\boldsymbol{\theta}$, for a particular choice of recurrent function and loss function.

Now we show why BPTT arises in the context of training an RTRBM. For the RTRBM, the average loss function is the KL divergence between data and model distributions, $p(\{\mathbf{y}\}_0^T)$ and $q(\{\mathbf{y}\}_0^T;\boldsymbol{\theta})$, respectively. The latter, however, can be re-expressed using Eq. 12 and the graph for the RTRBM, Fig. 3b:

$$\begin{aligned}q(\{\mathbf{y}\}_0^T;\boldsymbol{\theta}) &= \prod_{t=0}^{T}\int_{\mathbf{u}_{t-1}} q(\mathbf{y}_t|\mathbf{u}_{t-1};\boldsymbol{\theta})p(\mathbf{u}_{t-1}|\{\mathbf{y}\}_0^{t-1};\boldsymbol{\theta})\,\mathrm{d}\mathbf{u}_{t-1} \\ &= \prod_{t=0}^{T}\int_{\mathbf{u}_{t-1}} q(\mathbf{y}_t|\mathbf{u}_{t-1};\boldsymbol{\theta})\delta(\mathbf{u}_{t-1}-\bar{\mathbf{z}}_{t-1}(\{\mathbf{y}\}_0^{t-1},\boldsymbol{\theta}))\,\mathrm{d}\mathbf{u}_{t-1} \\ &= \prod_{t=0}^{T} q(\mathbf{y}_t|\bar{\mathbf{z}}_{t-1}(\{\mathbf{y}\}_0^{t-1},\boldsymbol{\theta});\boldsymbol{\theta}).\end{aligned}$$

The first line follows from the graph; the second from Eq. 12. Here, $\bar{\mathbf{z}}_t$ is the posterior mean, and therefore computable with a deterministic "upward" pass through the model,

$$\bar{\mathbf{z}}_t = f(W_{uz}\bar{\mathbf{z}}_{t-1} + W_{yz}\mathbf{y}_t + \mathbf{b}_z), \tag{25}$$

where the function f, acting element-wise on its vector argument, is determined by the exponential family to which the posterior distribution belongs (logistic for Bernoulli, exponential for Poisson, etc.). Now writing the KL divergence in terms of our new expression

for the model distribution, we find that

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'} \mathrm{D}_{\mathrm{KL}}\left\{p(\{\mathbf{y}\}_0^T)\middle|\middle|q(\{\mathbf{y}\}_0^T;\boldsymbol{\theta})\right\} &= -\frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'}\left\langle \log q(\{\mathbf{y}\}_0^T;\boldsymbol{\theta})\right\rangle_{p(\{\mathbf{y}\}_0^T)} \\
&= -\frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'}\left\langle \log \prod_{t=0}^{T} q(\mathbf{y}_t|\bar{\mathbf{z}}_{t-1}(\{\mathbf{y}\}_0^{t-1},\boldsymbol{\theta});\boldsymbol{\theta})\right\rangle_{p(\{\mathbf{y}\}_0^T)} \\
&= -\frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'}\left\langle \sum_{t=0}^{T} \log q(\mathbf{y}_t|\bar{\mathbf{z}}_{t-1}(\{\mathbf{y}\}_0^{t-1},\boldsymbol{\theta});\boldsymbol{\theta})\right\rangle_{p(\{\mathbf{y}\}_0^T)} \\
&= \frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'}\left\langle \sum_{t=0}^{T} L_t(\bar{\mathbf{z}}_{t-1}(\boldsymbol{\theta}),\boldsymbol{\theta})\right\rangle_{p(\{\mathbf{y}\}_0^T)} = \frac{\mathrm{d}}{\mathrm{d}\boldsymbol{\theta}'}\left\langle L\right\rangle_{p(\{\mathbf{y}\}_0^T)},
\end{aligned}$$

where we have defined

$$\begin{aligned}
L_t(\bar{\mathbf{z}}_{t-1}(\boldsymbol{\theta}),\boldsymbol{\theta}) &:= -\log q(\mathbf{y}_t|\bar{\mathbf{z}}_{t-1};\boldsymbol{\theta}), \\
L(\bar{\mathbf{z}}_0(\boldsymbol{\theta}),\ldots,\bar{\mathbf{z}}_{T-1}(\boldsymbol{\theta})) &:= \sum_{t=0}^{T} L_t(\bar{\mathbf{z}}_{t-1}(\boldsymbol{\theta}),\boldsymbol{\theta}).
\end{aligned} \tag{26}$$

So improving (decreasing) this loss is equivalent to making the RTRBM a good generative model for the data.

Comparing Eqs. 25 and 26 with Eqs. 21 and 22, we see that the RTRBM has the form of a recurrent neural network, with the posterior means playing the role of the hidden/output units, and the loss function depending on the parameters directly, as well as indirectly through its dependence on these "outputs." Applying BPTT, Eqs. 23 and 24, to this recurrent function and this loss means working out $\partial L/\partial \bar{\mathbf{z}}_t{}'$, $\partial L/\partial \boldsymbol{\theta}'$, $\partial \bar{\mathbf{z}}_{t+1}/\partial \bar{\mathbf{z}}_t{}'$, and $\partial \bar{\mathbf{z}}_t/\partial \boldsymbol{\theta}'$.

We start with the partial derivatives of the loss. First, from Eq. 26, they decouple over time:

$$\frac{\partial L}{\partial \boldsymbol{\theta}'} = \sum_{t=0}^{T} \frac{\partial L_t}{\partial \boldsymbol{\theta}'}, \qquad \frac{\partial L}{\partial \bar{\mathbf{z}}_t{}'} = \frac{\partial L_{t+1}}{\partial \bar{\mathbf{z}}_t{}'}. \tag{27}$$

Next, we recall that the distribution in Eq. 26 is the marginal over the visible units of an EFH (or RBM). (Recall from Eq. 15 in the main text that each RBM or EFH in the RTRBM is defined via a conditional distribution, $q(\mathbf{y}_{t+1},\mathbf{z}_{t+1}|\bar{\mathbf{z}}_t;\boldsymbol{\theta})$.) Hence, the partial derivatives with respect to the parameters, $\partial L_t/\partial W_{uz}$, $\partial L_t/\partial W_{yz}$, and $\partial L_t/\partial \mathbf{b}_z$, are precisely the terms that are normally computed during unsupervised learning of the inputs, $\mathbf{y}_{t+1}$, in a standard EFH. In the RTRBM, these gradients can be approximated with contrastive divergence (CD).

The gradient of the loss with respect to the hiddens/outputs—the quantity in the second equality in Eq. 27—turns out also to be computable with the standard unsupervised method for EFHs. Beginning with the well known formula for the gradient of the log of the

Boltzmann distribution,

$$\begin{aligned}
\frac{\partial L_{t+1}}{\partial \bar{\mathbf{z}}_t{}'} &= \mathbb{E}_{q(\mathbf{z}_{t+1},\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t)}\left[\frac{\partial H}{\partial \bar{\mathbf{z}}_t{}'}\middle|\bar{\mathbf{z}}_t\right] - \mathbb{E}_{q(\mathbf{z}_{t+1}|\mathbf{y}_{t+1},\bar{\mathbf{z}}_t)}\left[\frac{\partial H}{\partial \bar{\mathbf{z}}_t{}'}\middle|\mathbf{y}_{t+1},\bar{\mathbf{z}}_t\right], \\
H &= \mathbf{Z}_{t+1}{}'W_{yz}\mathbf{Y}_{t+1} + \mathbf{b}_y'\mathbf{Y}_{t+1} + (W_{uz}\bar{\mathbf{z}}_t + \mathbf{b}_z)'\mathbf{Z}_{t+1}. \\
\implies \frac{\partial H}{\partial \bar{\mathbf{z}}_t{}'} &= \mathbf{Z}_{t+1}{}'W_{uz} \\
&= \frac{\partial H}{\partial \mathbf{b}_z{}'}W_{uz}. \\
\implies \frac{\partial L_{t+1}}{\partial \bar{\mathbf{z}}_t{}'} &= \mathbb{E}_{q(\mathbf{z}_{t+1},\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t)}\left[\frac{\partial H}{\partial \mathbf{b}_z{}'}\middle|\bar{\mathbf{z}}_t\right] - \mathbb{E}_{q(\mathbf{z}_{t+1}|\mathbf{y}_{t+1},\bar{\mathbf{z}}_t)}\left[\frac{\partial H}{\partial \mathbf{b}_z{}'}\middle|\mathbf{y}_{t+1},\bar{\mathbf{z}}_t\right]W_{uz} \\
&= \frac{\partial L_{t+1}}{\partial \mathbf{b}_z{}'}W_{uz}.
\end{aligned} \tag{28}$$

Now we consider the derivatives of the hidden/output units. From Eq. 25, we see that

$$\frac{\partial \bar{\mathbf{z}}_{t+1}}{\partial \bar{\mathbf{z}}_t{}'} = J_{t+1}W_{uz}, \tag{29}$$

where J_{t+1} is the Jacobian of f, evaluated at $W_{uz}\bar{\mathbf{z}}_t + W_{yz}\mathbf{y}_{t+1} + \mathbf{b}_z$. The derivatives with respect to the parameters, for their part, are most felicitously expressed in connection with the row vector $\mathrm{d}L/\,\mathrm{d}\bar{\mathbf{z}}_t{}'$ (cf. Eq. 23):

$$\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}\frac{\partial \bar{\mathbf{z}}_t}{\partial W_{uz}} = J_t'\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t}\bar{\mathbf{z}}_{t-1}{}', \qquad \frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}\frac{\partial \bar{\mathbf{z}}_t}{\partial W_{yz}} = J_t'\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t}\mathbf{y}_t{}', \qquad \frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}\frac{\partial \bar{\mathbf{z}}_t}{\partial \mathbf{b}_z{}'} = \frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}J_t. \tag{30}$$

Filling Eqs. 27, 28, 29, and 30 into the equations for backprop-through-time, Eqs. 23 and 24, yields

$$\begin{aligned}
\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'} &= \frac{\partial L_{t+1}}{\partial \mathbf{b}_z{}'}W_{uz} + \frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_{t+1}{}'}J_{t+1}W_{uz}; \\
\frac{\mathrm{d}L}{\mathrm{d}W_{uz}} &= \left\langle \sum_{t=0}^{T}\frac{\partial L_t}{\partial W_{uz}} + \sum_{t=0}^{T}J_t'\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t}\bar{\mathbf{z}}_{t-1}{}' \right\rangle_{p(\{\mathbf{y}\}_0^T)}, \\
\frac{\mathrm{d}L}{\mathrm{d}W_{yz}} &= \left\langle \sum_{t=0}^{T}\frac{\partial L_t}{\partial W_{yz}} + \sum_{t=0}^{T}J_t'\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t}\mathbf{y}_t{}' \right\rangle_{p(\{\mathbf{y}\}_0^T)}, \\
\frac{\mathrm{d}L}{\mathrm{d}\mathbf{b}_z{}'} &= \left\langle \sum_{t=0}^{T}\frac{\partial L_t}{\partial \mathbf{b}_z{}'} + \sum_{t=0}^{T}\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}J_t \right\rangle_{p(\{\mathbf{y}\}_0^T)}.
\end{aligned} \tag{31}$$

These expressions can be economized if we work with

$$\mathbf{g}_t' := -\frac{\mathrm{d}L}{\mathrm{d}\bar{\mathbf{z}}_t{}'}J_t,$$

in which case, applying also Eq. 26, the complete RTRBM update equations are

$$\begin{aligned}
\mathbf{g}'_t &= \left(\frac{\partial \log q(\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t;\boldsymbol{\theta})}{\partial \mathbf{b}_z{}'} + \mathbf{g}'_{t+1}\right) W_{uz} J_t; \\
\Delta W_{uz} \propto -\frac{\mathrm{d}L}{\mathrm{d}W_{uz}} &= \left\langle \sum_{t=0}^{T} \left(\frac{\partial \log q(\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t;\boldsymbol{\theta})}{\partial W_{uz}} + \mathbf{g}_t \bar{\mathbf{z}}_{t-1}{}'\right)\right\rangle_{p(\{\mathbf{y}\}_0^T)}, \\
\Delta W_{yz} \propto -\frac{\mathrm{d}L}{\mathrm{d}W_{yz}} &= \left\langle \sum_{t=0}^{T} \left(\frac{\partial \log q(\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t;\boldsymbol{\theta})}{\partial W_{yz}} + \mathbf{g}_t \mathbf{y}_t{}'\right)\right\rangle_{p(\{\mathbf{y}\}_0^T)}, \\
\Delta \mathbf{b}_z \propto -\frac{\mathrm{d}L}{\mathrm{d}\mathbf{b}_z{}'} &= \left\langle \sum_{t=0}^{T} \left(\frac{\partial \log q(\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t;\boldsymbol{\theta})}{\partial \mathbf{b}_z{}'} + \mathbf{g}'_t\right)\right\rangle_{p(\{\mathbf{y}\}_0^T)}, \\
\Delta \mathbf{b}_y \propto -\frac{\mathrm{d}L}{\mathrm{d}\mathbf{b}_y{}'} &= \left\langle \sum_{t=0}^{T} \frac{\partial \log q(\mathbf{y}_{t+1}|\bar{\mathbf{z}}_t;\boldsymbol{\theta})}{\partial \mathbf{b}_y{}'}\right\rangle_{p(\{\mathbf{y}\}_0^T)}.
\end{aligned} \tag{32}$$

We have here also included the equation for updating the *visible* biases, which does not depend on the backwards recursion. In our experiments, each of the partial derivatives was approximated with contrastive divergence.

Appendix C. Details of the experiments

We provide here details about the generative model used to produce the "data distribution" on which the temporal EFH's were trained and tested in Section 4.2, and the testing procedure.

Details of the true generative model

The unobserved, linear dynamical system evolved according to

$$p(\mathbf{x}_{t+1}|\mathbf{x}_t) = \mathcal{N}(A\mathbf{x}_t, \Sigma_x), \qquad p(\mathbf{x}_0) = \mathcal{N}(\mathbf{0}, \Sigma_0). \tag{33}$$

To prevent these dynamics from being well approximated by a first-order dynamical model, we made two choices: (1) The state transition matrix was chosen to correspond to the discrete-time approximation to a(n undriven) damped harmonic oscillator, i.e., $m\ddot{x} + c\dot{x} + kx = 0$,

$$A = \begin{bmatrix} 1 & \delta t \\ -\frac{k}{m}\delta t & 1 - \frac{c}{m}\delta t \end{bmatrix},$$

with mass $m = 5$ kg, viscous damping $c = 0.25$ Ns/m, ideal-spring stiffness $k = 3$ N/m, and sampling interval $\delta t = 0.05s$. This makes the system stable and oscillatory (which first-order models cannot be). (2) Σ_x was chosen so that the Hankel singular values (see Scherpen, 2015) for the system, with output matrix $C = \begin{bmatrix} 1 & 0 \end{bmatrix}$ and input matrix set by Σ_x, were within one order of magnitude of each other; that is, ensuring that the transfer function from noise to position had roughly equal power in all modes. This was achieved with $\Sigma_x = \mathrm{diag}(\begin{bmatrix} 5\text{E-}7, & 5\text{E-}5 \end{bmatrix})$.

The current (time t) position was noisily encoded in the spike counts of a population of fifteen neurons, whose Gaussian-shaped tuning curves (f_i) evenly tile the space. Spike counts were drawn from (conditionally) independent Poisson distributions:

$$p(\mathbf{y}_t|\mathbf{x}_t, g_t) = \prod_{i=1}^{15} \text{Pois}[y_{i,t}|g_t f_i(x_t)]. \tag{34}$$

Here g_t is the "population gain," a scaling factor for the mean spike counts (Ma et al., 2006). Because the signal-to-noise ratio increases with mean for Poisson random variables, it essentially scales (linearly) the reliability of the population. Thus, in order to model instant-to-instant changes in sensory reliability, the gain was chosen independently and uniformly:

$$p(g_t) = \mathcal{U}\big(6.4, 9.6\big). \tag{35}$$

The gains are unobserved. The joint distribution of the states and their observations and the gains is the product of Eq. 33–Eq. 35, multiplied across all time. This distribution is consistent with the generic dynamical generative model in the first two rows of Fig. 1a.

In accordance with the broad tuning of higher sensory areas, the "standard deviation," σ_{tc}, of the tuning curves,

$$f_i(x) = \exp\left\{ -\frac{(x-\xi_i)^2}{2\sigma_{\text{tc}}^2} \right\},$$

was chosen so that the full-width at half maximum is one-sixth of the total space, for all preferred stimuli ξ_i.

This system is input-output stable (eigenvalues of the state-transition matrix are within the unit circle), but trajectories are nevertheless unbounded, since the input noise is unbounded (normally distributed). Thus, the (rare) stimuli that leave the space spanned by the tuning curves were simply "wrapped" onto the opposite side.

Although we have written the initial-state distribution as Gaussian, the initial position was in fact drawn from a uniform distribution across the space spanned by the Poisson "neurons." For EM learning, this was treated as an infinite-covariance Gaussian centered in the middle of joint space. The initial velocity was normally distributed very tightly about zero, so all movement of the oscillator was essentially due to initial displacement.

Details of the testing procedure

Models were tested on 40000 vectors, consisting of 40 trajectories of 1000 time steps apiece. For each trajectory, the vector of hidden means was computed from the current (time t) inputs according to Eq. 12. Then the position, X_t, was decoded from the hidden-unit means, $\bar{\mathbf{z}}_t$, using a scheme we have proposed previously for similar but static models (Makin et al., 2013). First, these activities were passed back down through the harmonium (i.e., through the weight matrix and element-wise nonlinearity) into the 15 "sensory" (i.e., the Poisson) units; and then the center of mass of these units was computed. Since the center of mass is the maximum-likelihood estimator for such a population (Makin et al., 2013), this is sensible, although not provably optimal, so the model performances presented in Fig. 4 is in some sense a lower bound. However, since the same decoding technique was used on all models, the comparisons between them are valid.

The same data were used for all tests in Fig. 4, so all variation across models and benchmarks results from the training alone.

References

Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37 (6A):3099–3132, dec 2009. ISSN 0090-5364. URL `http://projecteuclid.org/euclid.aos/1250515381`.

Zoubin Ghahramani and Geoffrey E. Hinton. Parameter Estimation for Linear Dynamical Systems. Technical report, University of Toronto, 1996.

Geoffrey E. Hinton. Training Products of Experts by Minimizing Contrastive Divergence. *Neural Computation*, 14:1771–1800, 2002.

Geoffrey E. Hinton. A Practical Guide to Training Restricted Boltzmann Machines. Technical report, University of Toronto, Toronto, 2010.

Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, 18:1527–1554, 2006. URL `http://www.mitpressjournals.org/doi/abs/10.1162/neco.2006.18.7.1527`.

Wei Ji Ma, Jeffrey M. Beck, Peter E. Latham, and Alexandre Pouget. Bayesian Inference with Probabilistic Population Codes. *Nature Neuroscience*, 9:1423–1438, 2006.

Joseph G. Makin and Philip N. Sabes. Sensory Integration and Density Estimation. *Advances in Neural Information Processing Systems 27: Proceedings of the 2014 Conference*, pages 1–9, 2015.

Joseph G. Makin, Matthew R. Fellows, and Philip N. Sabes. Learning Multisensory Integration and Coordinate Transformation via Density Estimation. *PLoS Computational Biology*, 9(4):1–17, 2013.

Joseph G. Makin, Benjamin K. Dichter, and Philip N. Sabes. Learning to Estimate Dynamical State with Probabilistic Population Codes. *PLOS Computational Biology PLoS Comput Biol*, 11(11):1–28, 2015. ISSN 1553-7358.

Kevin P. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, University of California, Berkeley, 2002.

J.M.A. Scherpen. *The Control Handbook*, chapter 4: Balanced Realizations, Model Order Reduction, and the Hankel Operator, pages 4–1–4–24. CRC Press, 2nd edition, 2015.

Ilya Sutskever. *Training Recurrent Neural Networks*. PhD thesis, University of Toronto, 2013.

Ilya Sutskever and Geoffrey E. Hinton. Learning Multilevel Distributed Representations for High-Dimensional Sequences. In *AISTATS*, pages 1–8, 2007.

Ilya Sutskever, Geoffrey E. Hinton, and Graham Taylor. The Recurrent Temporal Restricted Boltzmann Machine. In *Advances in Neural Information Processing Systems 21: Proceedings of the 2008 Conference*, pages 1–8, 2009.

D.M. Titterington, A.F.M. Smith, and U.E. Makov. *Statistical Analysis of Finite Mixture Distributions*. Wiley, 1985.

Max Welling, Michal Rosen-Zvi, and Geoffrey E. Hinton. Exponential Family Harmoniums with an Application to Information Retrieval. In *Advances in Neural Information Processing Systems 17: Proceedings of the 2004 Conference*, pages 1481–1488., 2005.